

Abstract of the Disclosure

The location of an object in a physical scene is identified with a stereo camera. A virtual surface is identified in the physical scene, and an approximate disparity set is constructed for the virtual surface. A main and a reference image of the scene are acquired by the stereo camera. The reference image is warped according to the disparity set, and subtracted from the main image to determine a set of depth residual values. Pixels having a substantially non-zero residual are identified with a surface of the object not coincident with the virtual surface.